

PATENT APPLICATION

ENHANCED VIRTUAL MACHINE INSTRUCTIONS

Inventors: 1. Stepan SOKOLOV
34832 Dorado Common
Fremont, CA 94555
Citizen of Ukraine

2. David WALLMAN
777 S. Mathilda Ave. #266
Sunnyvale, CA 94087
Citizen of Israel

Assignee: Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303

BEYER WEAVER & THOMAS, LLP
P.O. Box 778
Berkeley, CA 94704-0778
Telephone (650) 961-8300

ENHANCED VIRTUAL MACHINE INSTRUCTIONS

CROSS-REFERENCE TO RELATED APPLICATIONS

5

This application is related to U.S. Patent Application No. _____
(Att.Dkt.No. SUN1P811/P5512), entitled "REDUCED INSTRUCTION SET FOR
JAVA VIRTUAL MACHINES", filed on an even date, and hereby incorporated
herein by reference for all purposes.

10 This application is related to U.S. Patent Application No. 09/703,361
(Att.Dkt.No. SUN1P809/P5500), entitled "IMPROVED FRAMEWORKS FOR
INVOKING METHODS IN VIRTUAL MACHINES", which is hereby
incorporated herein by reference for all purposes.

15 This application is related to U.S. Patent Application No. 09/703,356
(Att.Dkt.No. SUN1P810/P5510), entitled "IMPROVED METHODS AND
APPARTUS FOR NUMERIC CONSTANT VALUE INLINING IN VIRTUAL
MACHINES", which is hereby incorporated herein by reference for all
purposes.

20 This application is related to U.S. Patent Application No. 09/703,449
(Att.Dkt.No. SUN1P814/P5417), entitled "IMPROVED FRAMEWORKS FOR
LOADING AND EXECUTION OF OBJECT-BASED PROGRAMS", which is
hereby incorporated herein by reference for all purposes.

BACKGROUND OF THE INVENTION

5 The present invention relates generally to object-based high level programming environments, and more particularly, to virtual machine instruction sets suitable for execution in virtual machines operating in portable, platform independent programming environments

10 Recently, the Java™ programming environment has become quite popular. The Java™ programming language is a language that is designed to be portable enough to be executed on a wide range of computers ranging from small devices (e.g., pagers, cell phones and smart cards) up to supercomputers. Computer programs written in the Java programming language (and other languages) may be compiled into Java Bytecode instructions that are suitable for execution by a Java virtual machine implementation.

15 The Java virtual machine is commonly implemented in software by means of an interpreter for the Java virtual machine instruction set but, in general, may be software, hardware, or both. A particular Java virtual machine implementation and corresponding support libraries, together constitute a Java™ runtime environment.

20 Computer programs in the Java programming language are arranged in one or more classes or interfaces (referred to herein jointly as classes or class files). Such programs are generally platform, i.e., hardware and operating system, independent. As such, these computer programs may be executed without modification, on any computer that is able to run an implementation 25 of the Java™ runtime environment. A class written in the Java programming language is compiled to a particular binary format called the "class file format" that includes Java virtual machine instructions for the methods of a single class. In addition to the Java virtual machine instructions for the methods of a class, the class file format includes a significant amount of ancillary

information that is associated with the class. The class file format (as well as the general operation of the Java virtual machine) is described in some detail in The Java Virtual Machine Specification by Tim Lindholm and Frank Yellin (ISBN 0-201-31006-6), which is hereby incorporated herein by reference.

5 Conventional virtual machines interpreter decodes and executes the Java Bytecode instructions, one instruction at a time during execution, e.g., "at runtime." To execute a Java instruction, typically, several operations have to be performed to obtain the information that is necessary to execute the Java instruction. For example, to invoke a method referenced by a Java bytecode,
10 the virtual machine must perform several operations to access the Constant Pool simply to identify the information necessary to locate and access the invoked method.

As described in The Java Virtual Machine Specification, one of the structures of a standard class file is known as the "Constant Pool." The
15 Constant Pool is a data structure that has several uses. One of the uses of the Constant Pool that is relevant to the present invention is that the Constant Pool contains the information that is needed to resolve various Java Instructions. To illustrate, Fig. 1 depicts a conventional computing environment 100 including a stream of Java Bytecodes 102, a constant pool 104 and an execution stack
20 106. The stream of Java Bytecodes 102 represents a series of bytes in a stream where one or more bytes can represent a Java Bytecode instruction. For example, a byte 108 can represent a Ldc (load constant on the execution stack) Bytecode command 108. Accordingly, the bytes 110 and 112 represent the parameters for the Ldc Bytecode command 108. In this case, these bytes
25 respectively represent a CP-IndexA 100 and CP-IndexB 112 that collectively represent the index to appropriate constant value in the constant pool 104. For example, bytes C1, C2, C3 and C4 of the constant pool 104 can collectively represent the appropriate 4 byte (one word) constant C that is to be loaded to the top of the execution stack 106. It should be noted that Ldc

Bytecode command 108 and its parameters represented by bytes 110 and 112 are collectively referred to herein as a Java Bytecode instruction.

In order to execute the Java Bytecode Ldc Instruction 108, at run time, an index to the Constant Pool 104 is constructed from the CP-IndexA and CP-IndexA. Once an index to the Constant Pool has been determined, the appropriate structures in the Constant Pool have to be accessed so that the appropriate constant value can be determined. Accordingly, the Java Bytecode Ldc instruction can be executed only after performing several operations at run time. As can be appreciated from the example above, the execution of a relatively simple instruction such as loading a constant value can take a significant amount of run time. Hence, this conventional technique is an inefficient approach that may result in significantly longer execution times.

Another problem is that the conventional Java Bytecode instruction set has more than 220 instructions. Moreover, there is a significant amount of redundancy between some instructions in the conventional Java Bytecode instruction set. For example, there are different Java Bytecode instructions for storing (or pushing) integer local variables on the execution stack (e.g., iLoad), and storing (or pushing) a pointer local variable on the execution stack (e.g., aLoad). However, the operations performed by these instructions are virtually identical, namely, storing (or pushing) 4 byte values (a word) on the execution stack. There is also a significant amount of overlap between some instructions of the conventional Java Bytecode instruction set. For example, there are 5 different Java Bytecode instructions for pushing one byte integer values on the execution stack (i.e., iconst_1, iconst_2, iconst_3, iconst_4 and iconst_5). However, these operations virtually perform the same operations, namely, pushing a constant one byte integer value on the execution stack.

As noted above, the Java Bytecode instruction set has more than 220 instructions. This means that conventionally nearly all of the 256 (2^8) allowable Bytecode values have to be assigned to Java instructions (commands

or opcodes). As a result, Java interpreters are needlessly complex since they need to recognize a relatively large number of Java instructions and possibly implement various mechanisms for executing many instructions. Thus, the conventional Java Bytecode instruction set is not a very desirable solution for 5 systems with limited resources (e.g., embedded systems)

Accordingly, there is a need for alternative instructions suitable for execution in virtual machines.

10

SUMMARY OF THE INVENTION

To achieve the foregoing and other objects of the invention, techniques for implementing virtual machine instructions suitable for execution in virtual machines are disclosed. The inventive virtual machine instructions can 15 effectively represent the complete set of operations performed by the conventional Java Bytecode instruction set. Moreover, the operations performed by conventional instructions can be performed by relatively fewer inventive virtual machine instructions. Furthermore, the inventive virtual machine instructions can be used to perform operations that cannot readily be 20 performed by conventional Java Bytecode instructions. Thus, a more elegant yet robust virtual machine instruction set can be implemented. This in turn allows implementation of relatively simpler interpreters as well as allowing alternative uses of the limited 256 (2^8) Bytecode representation (e.g., a macro representing a set of commands). As a result, the performance of virtual 25 machines, especially, those operating in systems with limited resources, can be improved.

The invention can be implemented in numerous ways, including a system, an apparatus, a method or a computer readable medium. Several embodiments of the invention are discussed below.

As a set of virtual machine instructions suitable for execution in a virtual machine to load constant values on an execution stack, one embodiment of the invention provides instructions representing a number of 5 corresponding Java Bytecode executable instructions that are also suitable for execution in the virtual machine to load constant values on an execution stack. The set of the virtual machine instructions consists of a number of virtual machine instructions that is less than the number of the corresponding Java Bytecode executable instructions. In addition, every one of the 10 corresponding Java Bytecode executable instructions can be represented by at least one of the virtual machine instructions in the virtual machine instruction set.

As a set of virtual machine instructions suitable for execution in a virtual machine to store local variables onto an execution stack, one 15 embodiment of the invention provides virtual machine instructions representing a number of corresponding Java Bytecode executable instructions that are also suitable for execution in the virtual machine to store local variables onto an execution stack. The set of virtual machine instructions consists of a number of virtual machine instructions that is less than the 20 number of the corresponding Java Bytecode executable instructions. In addition, every one of the corresponding Java Bytecode executable instructions can be represented by at least one of the virtual machine instructions in the virtual machine instruction set.

As a virtual machine instruction suitable for execution in a virtual 25 machine to load values from arrays on an execution stack, one embodiment of the invention provides a virtual machine instruction representing two or more Java Bytecode executable instructions that are also suitable for loading values from arrays on the execution stack.

As a virtual machine instruction suitable for execution in a virtual 30 machine to store values located on an execution stack into arrays, one

embodiment of the invention provides a virtual machine instruction representing two or more Java Bytecode executable instructions that are also suitable for storing values located on an execution stack into an array.

One embodiment of the invention provides a virtual machine

5 instruction suitable for execution in a virtual machine to duplicate values stored in an execution stack on top of the execution stack. The virtual machine instruction represents two or more Java Bytecode executable instructions that are also suitable for duplicating values stored in the execution stack on top of the execution stack.

10 Another embodiment of the invention provides a virtual machine instruction suitable for execution in a virtual machine to duplicate values stored in an execution stack on top of the execution stack, wherein the virtual machine instruction has a parameter associated with it to indicate which value stored in the execution stack should be duplicated on the top of the stack.

15 As a virtual machine instruction suitable for execution in a virtual machine, one embodiment of the invention operates to return values by placing them on top of an execution stack. The virtual machine instruction represents two or more Java Bytecode executable instructions that are also suitable for returning values by placing them on top of the execution stack.

20 Other aspects and advantages of the invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating by way of example the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein 5 like reference numerals designate like structural elements, and in which:

Fig. 1 depicts a conventional computing environment including a stream of Java Bytecodes, a constant pool, and an execution stack.

Fig. 2A is a block diagram representation of a computing environment including a Java Bytecode instruction translator in accordance with one 10 embodiment of the invention.

Fig. 2B illustrates a mapping that can be performed by the Java Bytecode instruction translator of Fig. 2A in accordance with one embodiment of the invention.

Fig. 3 illustrates an internal representation of Java instructions in 15 accordance with one embodiment of the invention.

Fig. 4A illustrates an internal representation of a set of Java Load Constant Bytecode instructions in accordance with one embodiment of the invention.

Fig. 4B illustrates a set of conventional Java Bytecode instructions that 20 can be represented by an inventive Push command.

Fig. 4C illustrates an internal representation of a set of conventional Java Load Constant Bytecode instructions in accordance with another embodiment of the invention.

Fig. 4D illustrates a set of conventional Java Bytecode instructions that 25 can be represented by a single PushL command in accordance with one embodiment of the invention.

Fig. 4E illustrates an internal representation of a set of Java Load Constant Bytecode instructions in accordance with yet another embodiment of the invention.

Fig. 4F illustrates a set of Java Bytecode instructions that can be 5 represented by a single PushB command in accordance with one embodiment of the invention.

Fig. 5A illustrates an internal representation of a set of Java Load from a local variable instructions in accordance with another embodiment of the invention.

10 Fig. 5B illustrates a set of Java Bytecode instructions for loading 4 byte local variables that can be represented by an inventive Load command in accordance with one embodiment of the invention.

Fig. 5C illustrates a set of Java Bytecode instructions for loading 8 byte local variables in accordance with one embodiment of the invention.

15 Fig. 6A illustrates a computing environment including an ALoad (load from array) virtual machine instruction in accordance with one embodiment of the invention.

20 Fig. 6B illustrates a set of conventional Java Bytecode instructions for loading arrays that can be represented by a single inventive virtual machine instruction in accordance with one embodiment of the invention.

Fig. 6C illustrates a computing environment including an AStore (store into array) virtual machine instruction in accordance with one embodiment of the invention.

25 Fig. 6D and 6E illustrate a set of conventional Java Bytecode instructions for storing arrays that can be represented by a single inventive virtual machine instruction.

Fig. 6F and 6G illustrate some Java conventional Bytecode instructions for performing conditional flow operations which can be represented by two

inventive virtual machine instructions in accordance with one embodiment of the invention.

Fig. 7A illustrates a computing environment including an internal representation of a DUP instruction suitable for duplicating values on the stack
5 in accordance with one embodiment of the invention.

Figs. 7B and 7C illustrate various Java Bytecode instructions that can be represented by an inventive virtual machine instruction in accordance with one embodiment of the invention.

Figs. 8A and 8B illustrate mapping of Java Bytecode return instructions
10 to virtual machine instructions provided in accordance with one embodiment of the invention.

Fig. 9 illustrates a mapping of Java Bytecode instantiation instructions to the virtual machine instructions provided in accordance with one embodiment of the invention

15

DETAILED DESCRIPTION OF THE INVENTION

As described in the background section, the Java programming
20 environment has enjoyed widespread success. Therefore, there are continuing efforts to extend the breadth of Java compatible devices and to improve the performance of such devices. One of the most significant factors influencing the performance of Java based programs on a particular platform is the performance of the underlying virtual machine. Accordingly, there have been
25 extensive efforts by a number of entities to improve performance in Java compliant virtual machines.

To achieve this and other objects of the invention, techniques for implementing virtual machine instructions suitable for execution in virtual

machines are disclosed. The inventive virtual machine instructions can effectively represent the complete set of operations performed by the conventional Java Bytecode instruction set. Moreover, the operations performed by conventional instructions can be performed by relatively fewer 5 inventive virtual machine instructions. Furthermore, the inventive virtual machine instructions can be used to perform operations that cannot readily be performed by the conventional Java Bytecode instructions. Thus, a more elegant yet robust virtual machine instruction set can be implemented. This, in turn, allows implementation of relatively simpler interpreters as well as 10 allowing for alternative uses of the limited 256 (2^8) Bytecode representation (e.g., a macro representing a set of commands). As a result, the performance of virtual machines, especially, those operating in systems with limited 15 resources, can be improved.

Embodiments of the invention are discussed below with reference to 15 Figs. 2 –9. However, those skilled in the art will readily appreciate that the detailed description given herein with respect to these figures is for explanatory purposes only as the invention extends beyond these limited 20 embodiments.

Fig. 2A is a block diagram representation of a computing environment 200 including a Java Bytecode instruction translator 202 in accordance with 20 one embodiment of the invention. The Java Bytecode instruction translator 202 operates to convert one or more bytes of a Java Bytecode stream 204, representing a Java Bytecode instruction 205 into a virtual machine instruction 206 containing one or more bytes. The Java Bytecode instruction 205 in the 25 Java Bytecode stream 204 can be, for example, a “Lcd” command 108 with its associated parameters 110 and 112, as described in Fig. 1.

Typically, one byte of the virtual machine instruction 206 is designated 30 to represent a virtual machine command (or opcode). In addition, one or more additional bytes may be associated with the virtual machine command (or opcode) to represent its parameters. As a result, one or more bytes of the

virtual machine instruction 206 can represent a Java Bytecode instruction having one or more bytes that collectively represent a Java Bytecode instruction, namely a command and possibly the parameters associated with the command (e.g., a one byte Java `iconst_1` Bytecode instruction, three bytes representing Java Bytecode `Ldc` instruction, etc).

As will be appreciated, the virtual machine instruction 206 can represent similar virtual machine operations that the Java Bytecode instruction 205 represents. In addition, the virtual machine instruction 206 can be loaded by a virtual machine instruction loader 208 into a virtual machine 210 as an internal representation 212. As will become apparent, the internal representation 212 can be used to significantly improve the performance of the virtual machine.

Furthermore, the Java Bytecode instruction translator 202 is capable of converting a set of Java Bytecode executable instructions into a more elegant set of instructions that is especially suitable for systems with limited resources. The operations performed by a conventional Bytecode instruction set can effectively be represented by fewer inventive virtual machine instructions. Accordingly, the number of the executable virtual machine instructions can be significantly less than the number of conventional Java Bytecode executable instructions needed to perform the same set of operations. In other words, two or more distinct conventional Java Bytecode executable instructions can effectively be mapped into an inventive virtual machine instruction.

To elaborate, Fig. 2B illustrates a mapping 250 that can be performed, for example, by the Java Bytecode instruction translator 202 in accordance with one embodiment of the invention. As illustrated in Fig. 2B, a set of conventional Java Bytecode executable instructions 252 can be mapped into a corresponding set of inventive virtual machine instructions 254. It should be noted that the set of Java Bytecode executable instructions 252 consists of M instructions, Bytecode Instructions (BC_1-BC_M). It should also be noted that each of the Byte code Instructions BC_1-BC_M represent a unique instruction in

the set of Java Bytecode executable instructions 252. As will be appreciated, the corresponding set of executable virtual machine instructions 254 consists of N instructions (AVM₁- AVM_N), a number that can be significantly less than M (the number of Java Bytecode executable instructions 252). Accordingly, 5 two or more Byte code Instructions of the Java Bytecode executable instructions 252 can be mapped into the same executable virtual machine instruction. For example, Bytecode Instructions BC_i, BC_j and BC_k can all be mapped into the same virtual machine executable instruction, namely, the instruction AVM₁. In addition, as will be described below, two or more 10 inventive virtual machine instructions from the set 254 can be combined to effectively represent a Java Bytecode instruction in the set 252.

As noted above, a virtual machine instruction, for example, the instruction AVM₁, can be loaded by a virtual machine instruction loader into a virtual machine as an internal representation that can be used to significantly 15 improve the performance of the virtual machine. Fig. 3 illustrates an internal representation 300 in accordance with one embodiment of the invention. The internal representation 300 can, for example, be implemented as a data structure embodied in a computer readable medium that is suitable for use by a virtual machine. As shown in Fig. 3, the internal representation 300 20 includes a pair of streams, namely, a code stream 302 and a data stream 304.

It should be noted that conventionally Java Bytecode instructions are internally represented as a single stream in the virtual machine. However, as shown in Fig. 3, the internal representation 300 includes a pair of streams, namely, a code stream 302 and a data stream 304. More details about 25 representing instructions as a pair of streams can be found in the U.S. Patent Application No. 09/703,449 (Att.Dkt.No. SUN1P814/P5417), entitled “IMPROVED FRAMEWORKS FOR LOADING AND EXECUTION OF OBJECT-BASED PROGRAMS”.

Each one of the entries in the code stream 302 and/or data stream 304 30 represent one or more bytes. The code stream 302 includes various virtual

machine commands (or instructions) 306, 308 and 310. The virtual machine commands (or instruction) 306 represents a virtual machine instruction that does not have any parameters associated with it. On the other hand, each of the virtual machine commands B and C have associated data parameters that 5 are represented in the data stream 304. More particularly, data B is the corresponding data parameter of the virtual machine command B, and data C1 and C2 are the data parameters associated with the command C.

It should be noted that the inventive virtual machine command B and data B represents one or more conventional Java Bytecodes which have been 10 converted, for example, by the Java Bytecode instruction translator 202 of Fig. 2A . Similarly, the virtual machine command C, data C1 and C2 collectively represent the one or more Java Bytecodes that have been converted into an inventive virtual machine instruction with its appropriate data parameters.

Fig. 4A illustrates an internal representation 400 of a set of Java Load 15 Constant Bytecode instructions in accordance with one embodiment of the invention. The internal representation 400 can, for example, be implemented as a data structure embodied in a computer readable medium that is suitable for use by a virtual machine.

In the described embodiment, each entry in the code stream 402 and 20 data stream 404 represents one byte. As such, the code stream 402 includes a one byte Push command 406, representing an inventive virtual machine command suitable for representation of one or more conventional Java Load Constant Bytecode instructions. The data stream 404 includes the data parameters associated with the Push command 406, namely, bytes A, B, C and 25 D. As will be appreciated, at execution time, the virtual machine can execute the Push command 406. Accordingly, the value represented by the bytes A, B, C and D in the data stream 404 can be pushed on the execution stack. In this way, the Push command 406 can effectively represent various Java Bytecode instructions that push values represented by 4 bytes (one word) on the 30 execution stack at run time. Fig. 4B illustrates a set of conventional Java

Bytecode instructions that can be represented by an inventive Push command (e.g., Push command 406).

Fig. 4C illustrates an internal representation 410 of a set of conventional Java Load Constant Bytecode instructions in accordance with another embodiment of the invention. Similar to the internal representation 400 of Fig. 4A, the internal representation 410 includes a pair of streams, namely, the code stream 402 and data stream 404, wherein each entry in the streams represents one byte. However, in Fig. 4B, the code stream 402 includes a one byte PushL command 412, representing another inventive virtual machine instruction suitable for representation of one or more Java Load Constant Bytecode instructions. It should be noted that the PushL command 412 has 8 bytes of data associated with it, namely, the bytes represented by A, B, C, D, E, F, G and H in the data stream 404. At execution time, the virtual machine can execute the PushL command 412 to push the value represented by the bytes A, B, C, D, E, F, G and H in the data stream 404, on the top of the execution stack. Accordingly, the PushL command 412 can effectively represent various Java Bytecode instructions that push 8 byte (two word) values on the execution stack at run time. Fig. 4D illustrates a set of conventional Java Bytecode instructions that can be represented by a single PushL command (e.g., PushL command 412) in accordance with one embodiment of the invention.

Fig. 4E illustrates an internal representation 420 of a set of Java Load Constant Bytecode instructions in accordance with yet another embodiment of the invention. Again, the internal representation 420 includes the code stream 402 and data stream 404, wherein each entry in the streams represents one byte. However, in Fig. 4E, the code stream 402 includes a one byte PushB command 422, representing yet another inventive virtual machine instruction suitable for representation of one or more Java Load Constant Bytecode instructions. It should be noted that the PushB command 422 has a one byte data parameter A associated with it. As shown in Fig. 4E, the data parameter

can be stored in the code stream 402. However, it should be noted that in accordance with other embodiment of the invention, the data parameter A can be stored in the data stream 404. In any case, the PushB command 422 can effectively represent various Java Bytecode instructions that push one byte 5 values on the execution stack at run time. Fig. 4F illustrates a set of Java Bytecode instructions that can be represented by a single PushB command (e.g., PushB command 422) in accordance with one embodiment of the invention.

Fig. 5A illustrates an internal representation 500 of a set of Java "Load 10 from a local variable" instructions in accordance with another embodiment of the invention. In the described embodiment, a code stream 502 of the internal representation 500 includes a Load command 506, representing an inventive virtual machine instruction suitable for representation of one or more Java "Load from a local variable" Bytecode instructions. It should be noted that the 15 Load command 506 has a one byte parameter associated with it, namely, an index_i 508 in the data stream 504. As will be appreciated, at run time, the Load command 506 can be executed by a virtual machine to load (or push) a local variable on top of the execution stack 520. By way of example, an offset₀ 522 can indicate the starting offset for the local variables stored on the 20 execution stack 520. Accordingly, an offset_i 524 identifies the position in the execution stack 520 which corresponds to the index_i 508 shown in Fig. 5A.

It should be noted that in the described embodiment, the Load 25 command 506 is used to load local variables as 4 bytes (one word). As a result, the value indicated by the 4 bytes A, B, C and D (starting at offset_i 524) is loaded on the top of the execution stack 520 when the Load command 506 is executed. In this manner, the Load command 506 and index_i 508 can be used to load (or Push) 4 byte local variables on top of the execution stack at 30 run time. As will be appreciated, the Load command 506 can effectively represent various conventional Java Bytecode instructions. Fig. 5B illustrates a set of Java Bytecode instructions for loading 4 byte local variables that can be

represented by an inventive Load command (e.g., Load command 412) in accordance with one embodiment of the invention.

It should be noted that the invention also provides for loading local variables that do not have values represented by 4 bytes. For example, Fig. 5C
5 illustrates a set of Java Bytecode instructions for loading 8 byte local variables in accordance with one embodiment of the invention. As will be appreciated, all of the Java Bytecode instructions listed in Fig. 5C can be represented by a single inventive virtual machine instruction (e.g., a LoadL command). The LoadL command can operate, for example, in a similar manner as discussed
10 above.

In addition, the invention provides for loading values from arrays into an execution stack. By way of example, Fig. 6A illustrates a computing environment 600 in accordance with one embodiment of the invention. The computing environment 600 includes an array 602 representative of a Java array stored in a portion of a memory of the computing environment 600. An execution stack 604 is also depicted in Fig. 6. As will be appreciated, an 15 inventive virtual machine instruction ALoad (array load) 605 can be utilized to facilitate loading of various values from the array 602 to the top of the execution stack 604.

20 During the execution of the virtual machine instruction ALoad 605, an array-reference 606 can be utilized (e.g., resolved) to determine the location of the array 602. In addition, an array-index 608 can be used to identify the appropriate offset of the array 602 and thereby indicate the appropriate value that is to be loaded from the array 602 on the execution stack 604. As will be
25 appreciated, the inventive virtual machine instruction ALoad can be used to load the appropriate values from various types of arrays (e.g., 1 byte, 2 bytes, 4 bytes, 8bytes arrays). To achieve this, a header 610 of the array 602 can be read to determine the arrays' type. Accordingly, based on the type of the array 602 as indicated by the header 610, the appropriate value that is to be loaded

from the array can be determined by using the array-index 608. This value can then be loaded onto the top of the execution stack 604.

Thus, the inventive virtual machine instruction ALoad can effectively represent various Java Bytecode instructions that are used to load values from an array. Fig. 6B illustrates a set of conventional Java Bytecode instructions for loading arrays that can be represented by a single inventive virtual machine instruction (e.g., ALoad) in accordance with one embodiment of the invention.

As will be appreciated, the invention also provides for virtual machine instructions used to store values into arrays. By way of example, Fig. 6C illustrates a computing environment 620 in accordance with one embodiment of the invention. An inventive AStore 622 (store into array) virtual machine instruction can be used to store various values from the execution stack 604 into different types of arrays in accordance with one embodiment of the invention. Again, the header 610 of the array 602 can be read to determine the array's type. Based on the array's type, the appropriate value (i.e., the appropriate number of bytes N on the execution stack 604 of Fig. 6B) can be determined. This value can then be stored in the array 602 by using the array-index 626. Thus, the inventive virtual machine instruction ALoad can effectively represent various Java Bytecode instructions that are used to store values into an array. Figs. 6D and 6E illustrate a set of conventional Java Bytecode instructions for storing arrays that can be represented by an inventive virtual machine instruction (e.g., Astore) in accordance with one embodiment of the invention.

Still further, two or more of the inventive virtual machine instructions can be combined to perform relatively more complicated operations in accordance with one embodiment of the invention. By way of example, the conditional flow control operation performed by the Java Bytecode instruction "Icmp" (compare two long values on the stack and based on the comparison push 0 or 1 on the stack) can effectively be performed by performing an

inventive virtual machine instruction LSUB (Long subdivision) followed by another inventive virtual machine instruction JMPEQ (Jump if equal). Fig. 6F and 6G illustrate some Java conventional Bytecode instructions for performing conditional flow operations which can be represented by two inventive virtual machine instructions in accordance with one embodiment of the invention.

The invention also provides for inventive operations that cannot be performed by Java Bytecode instructions. By way of example, an inventive virtual machine operation “DUP” is provided in accordance with one embodiment of the invention. The inventive virtual machine instruction DUP 10 allows values in various positions on the execution stack to be duplicated on the top of the execution stack. Fig. 7A illustrates a computing environment 700 including an internal representation 701 of a DUP instruction 702 suitable for duplicating values on the stack in accordance with one embodiment of the invention. The internal representation 701 includes a pair of streams, namely, 15 a code stream 402 and a data stream 404. In the described embodiment, each entry in the code stream 402 and data stream 404 represents one byte. The inventive virtual machine instruction DUP 702 is associated with a data parameter A in the code stream 402. Again, it should be noted that Data parameter A can be implemented in the data stream 404. In any case, the data 20 parameter A indicates which 4 byte value (word value) on an execution stack 704 should be duplicated on the top of the execution stack 704. The data parameter A can indicate, for example, an offset from the top of the execution stack 704. As shown in Fig. 7A, the data parameter A can be a reference to W_i , a word (4 byte) value on the execution stack. Accordingly, at execution 25 time, the virtual machine can execute the DUP command 702. As a result, the W_i word will be duplicated on the top of the stack. Thus, as will be appreciated, the inventive DUP virtual machine instructions can effectively replace various Java Byte instructions that operate to duplicate 4 byte values on top of the execution. Fig. 7B illustrates some of these Java Bytecode 30 instructions. Similarly, as illustrated in Fig. 7C, an inventive DUPL virtual machine can be provided to effectively replace various Java Bytecode

instructions that operate to duplicate 8 byte values (2 words) on top of the execution stack.

It should be noted that conventional Java Bytecode instructions only allow for duplication of values in certain positions on the execution stack (i.e., 5 dup, dup_x1 and dupx2 respectively allow duplication of W1, W2 and W3 on the stack). However, the inventive virtual machine instructions DUP and DUPL can be used to duplicate a much wider range of values on the execution stack (e.g., W4, Wi, WN, etc.)

Figs. 8A and 8B illustrate mapping of Java Bytecode "Return" 10 instructions to virtual machine instructions provided in accordance with one embodiment of the invention. As shown in Fig. 8A, various Java Bytecode instructions can be effectively mapped into a Return virtual machine instruction. As will be appreciated, the Return virtual machine instruction operates to put 4 byte values (one word) on the execution stack in a similar 15 manner as the virtual machine instructions for loading constants on the stack described above (e.g., iload). Fig. 8B illustrates a mapping of Java Bytecode return instructions to a "Lreturn" virtual machine instruction that can operate to put 8 byte values (two words) on the execution stack.

In a similar manner, Fig. 9 illustrates a mapping of Java Bytecode 20 instantiation instructions to the virtual machine instructions provided in accordance with one embodiment of the invention. Again, the four various Java Bytecode instructions can be effectively mapped into a virtual machine instruction (e.g., NEW). The virtual machine instruction NEW operates to instantiate objects and arrays of various types. In one embodiment, the 25 inventive virtual machine instruction NEW operates to determine the types of the objects or arrays based on the parameter value of the Bytecode instantiation instructions. As will be appreciated, the Bytecode instructions for instantiation are typically followed by a parameter value that indicates the type. Thus, the parameter value is readily available and can be used to allow

the NEW virtual machine instruction to instantiate the appropriate type at execution time.

Appendix A illustrates mapping of a set of conventional Java Bytecode instructions to one or more of the inventive virtual machine instructions listed
5 the in right column.

The many features and advantages of the present invention are apparent from the written description, and thus, it is intended by the appended claims to cover all such features and advantages of the invention. Further, since numerous modifications and changes will readily occur to those skilled in the
10 art, it is not desired to limit the invention to the exact construction and operation as illustrated and described. Hence, all suitable modifications and equivalents may be resorted to as falling within the scope of the invention.

What is claimed is: